

## IN THE CLAIMS:

The following listing of claims replaces all previous listings of claims in this application:

1. (Canceled)

2. (Previously Presented): A system as in claim 21, in which the memory reservation software module is a driver installed within each respective guest OS.

3. (Previously Presented): A system as in claim 2, in which there is a plurality of virtual computers operatively connected to the host system, further comprising:

a resource scheduler in the host system for allocating the hardware memory among the virtual computers;

for each virtual computer, a communications module for communicating a respective memory quantity request to each driver;

each driver, upon sensing the respective memory quantity request, being provided for causing the corresponding guest OS to reserve an amount of the guest physical memory corresponding to the memory quantity request.

4. (Previously Presented): A system as in claim 3, in which:

each guest OS includes computer-executable, memory reservation code for reserving specified amounts of the guest physical memory;

the driver is operatively connected to the memory reservation code for communicating the memory quantity request to the memory reservation code; and

the memory reservation code of each guest OS is native to the guest OS, all communication between the resource scheduler and the virtual computers taking place via the respective drivers, the resource scheduler thereby remaining transparent to the virtual computers.

5. (Cancelled)

6. (Previously Presented): A system as in claim 4, in which each virtual computer includes a virtual machine and a virtual machine monitor forming an interface between the resource scheduler and each respective virtual machine.

7. (Previously Presented): A system as in claim 4 in which:

the guest OS changes the amount of the guest physical memory allocated to applications and drivers loaded within and connected to the guest OS;

upon an increase in the memory quantity request issued by the resource scheduler for a specified one of the drivers, the corresponding specified guest OS reserves for the specified driver a corresponding quantity of guest physical memory, the driver thereby making the hardware memory corresponding to the reserved guest physical memory available for allocation by the host OS to other virtual computers; and

upon a decrease in the memory quantity request issued by the resource scheduler for the specified one of the drivers, the corresponding specified guest OS releases any prior reservation of a corresponding quantity of guest physical memory, thereby making the hardware memory corresponding to the released guest physical memory available for use by the specified virtual computer.

8. (Previously Presented): A system as in claim 21, in which the memory reservation software module is further provided for adapting a rate at which it reserves the guest physical memory via the guest OS to be no greater than a current maximum reservation change rate of the guest OS.

9. (Previously Presented): A system as in claim 21, in which the memory reservation software module is a user-level application loaded in the virtual and running on the guest OS.

10. (Cancelled)

11. (Cancelled)

12. (Currently Amended): A computer system comprising:

a host system, which includes a host operating system (OS) and a hardware memory forming a system resource; and

a plurality of virtual computers operatively connected to the host system, each of which includes at least one virtual processor, guest physical memory, and a guest OS operable to address the guest physical memory in a guest physical address space;

a resource scheduler in the host system for allocating the system resource among the virtual computers;

each guest OS being provided with a memory reservation software module for receiving a memory quantity request from the host system and for changing the allocation of the guest physical memory from within the respective guest OS, thereby changing the amount of the hardware memory available for arbitrary use by the host system;

in which:

the memory reservation software module is a driver installed within each respective guest OS;

each driver, upon sensing the respective memory quantity request, reserves, via the corresponding guest OS, an amount of the guest physical memory corresponding to the memory quantity request;

the driver is operatively connected to the resource scheduler for communicating the memory quantity request to the resource scheduler;

the memory reservation software module of each guest OS is native to the guest OS, all communication between the resource scheduler and the virtual computers taking place via the respective drivers, the resource scheduler thereby remaining transparent to the virtual computers;

the guest OS changes the amount of guest physical memory allocated to applications and drivers loaded within and connected to the guest OS;

upon an increase in the resource quantity request issued by the resource scheduler for a specified one of the drivers, the corresponding specified guest OS reserves for the specified driver a corresponding quantity of the guest physical memory at a controlled rate, said driver adjusting the controlled rate up or down between a

predetermined maximum rate and a predetermined minimum rate as said corresponding specified guest OS reserves the corresponding quantity of the guest physical memory for the specified driver at the controlled rate,

the driver thereby making the hardware memory corresponding to the reserved guest physical memory available for allocation by the host OS to other virtual computers; and

upon a decrease in the memory quantity request issued by the resource scheduler for the specified one of the drivers, the corresponding specified guest OS releases any prior reservation of a corresponding quantity of guest physical memory, thereby making the hardware memory corresponding to the released guest physical memory available for use by the specified virtual computer.

13. (Previously Presented): In a computer system that comprises a host system, which includes:

a host operating system (OS),

a hardware memory; and

at least one virtual computer, which includes a virtual processor, guest physical memory, and a guest OS operable to address and allocate the guest physical memory in a guest physical address space, the virtual computer being operatively connected as a guest system running on the host system,

a method comprising the step of changing the allocation of the guest physical memory from within the guest OS in response to a memory quantity request issued by the host system, thereby changing the amount of the hardware memory available for arbitrary use by the host system.

14. (Previously Presented): A method as in claim 13, further including the following steps:

communicating the memory quantity request from a resource scheduler in the host system to a driver located within the guest OS; and

reserving, via the corresponding guest OS, an amount of the guest physical memory corresponding to the memory quantity request.

15. (Previously Presented): A method as in claim 14, in which the step of reserving the amount of the guest physical memory is performed using a resource reservation mechanism that is native to the guest OS, all communication between the resource scheduler and the virtual computer taking place via the respective drivers, the resource scheduler thereby remaining transparent to the virtual computer.

16. (Previously Presented): A method as in claim 15, further comprising the following steps:

implementing each virtual computer as a virtual machine and a respective virtual machine monitor; and

providing communication between each virtual machine and the host system via the respective virtual machine monitor.

17. (Previously Presented): A method as in claim 16, further including the following steps:

allocating and deallocating guest physical memory from within the guest OS to applications and drivers loaded within and connected to the guest OS;

upon an increase in the memory quantity request issued by the resource scheduler for a specified one of the drivers, reserving, via the corresponding specified guest OS, a corresponding quantity of the guest physical memory, the hardware memory corresponding to the reserved guest physical memory thereby becoming available for allocation by the host OS to other virtual computers or to the host OS itself; and

upon a decrease in the memory quantity request for the specified one of the drivers, releasing any prior reservation of a corresponding quantity of the guest physical memory, thereby making the hardware memory corresponding to the released guest physical memory available for use by the specified virtual computer.

18. (Cancelled)

19. (Cancelled)

20. (Previously Presented): A method as in claim 17, further including the step of adapting a rate at which the guest physical memory is reserved via the guest OS to be no greater than a current maximum reservation change rate of the guest OS.

21. (Currently Amended): A computer system comprising:

a host system, which includes a host operating system (OS) and a hardware memory that is addressable in a hardware memory address space;

at least one virtual computer, each of which includes at least one virtual processor, guest physical memory, and a guest OS operable to address, allocate and deallocate the guest physical memory in a guest physical address space, and is operatively connected to the host system; and

a memory reservation software module located within the virtual computer for receiving a memory quantity request from the host system and for changing the allocation of the guest physical memory from within the respective guest OS according to the memory quantity request, thereby changing the amount of the hardware memory available for arbitrary use by the host system,

wherein the memory reservation software module, in response to the memory quantity request, changes the allocation of the guest physical memory from within the respective guest OS at a controlled rate, said controlled rate being adjusted up or down between a predetermined maximum rate and a predetermined minimum rate as the memory reservation software changes the allocation of the guest physical memory from within the respective guest OS,

said controlled rate being increased if said memory reservation software module is successful in changing the allocation of the guest physical memory from within the respective guest OS at the controlled rate, and

said controlled rate being decreased if said memory reservation software module is not successful in changing the allocation of the guest physical memory from within the respective guest OS at the controlled rate.

22. (Previously Presented): A computer system comprising:  
a host system, which includes a host operating system (OS);  
a plurality of cooperating hardware processors that are included in the host system and that are collocated in a single hardware platform;  
at least one guest system that is operatively connected to the host system;  
each guest system being provided with a processor reservation software module comprising computer-executable code for receiving from the host system a respective processor quantity request, which indicates a number of the plurality of processors to be reserved by each guest system, and for indicating to the guest system a change in the number of processors reserved for use by the respective guest system, thereby making the reserved processors available for arbitrary use by the host system.

23. (Previously Presented): A system as in claim 22, in which each guest system is a virtual computer, each of which includes at least one virtual processor, guest physical memory, and a guest OS.

24. (Previously Presented): A system as in claim 23, in which the processor reservation software module is a driver installed within each respective guest OS.

25. (Previously Presented): A system as in claim 23, further comprising:  
a resource scheduler in the host system for allocating the hardware processors among the virtual computers; and  
for each virtual computer, a communications module for communicating the respective processor quantity request to each driver.

26. (Previously Presented): A system as in claim 25, in which:  
the processor reservation software module of each guest OS is native to the guest OS, all communication between the resource scheduler and the virtual computers taking place via the respective drivers, the resource scheduler thereby remaining transparent to the virtual computers.

27. (Currently Amended): A computer system comprising:

a host system, which includes at least one hardware processor, a host operating system (OS) and hardware memory that is addressable in a hardware memory address space; and

a plurality of virtual computers, each of which comprises a body of computer-executable code stored in the hardware memory and includes at least one virtual processor, guest physical memory, and a guest OS operable to change allocation of the guest physical memory in a guest physical address space, and is operatively connected to the host system;

a memory reservation software module located within the virtual computer for receiving a memory quantity request from the host system and for changing the allocation of the guest physical memory from within the respective guest OS according to the memory quantity request, thereby changing the amount of the hardware memory available for arbitrary use by the host system;

the guest physical address spaces of all the virtual computers being mapped as portions of the hardware memory address space;

the virtual computers all being executable on the same hardware processor(s), which are collocated in a single hardware platform,

wherein the memory reservation software module, in response to the memory quantity request, changes the allocation of the guest physical memory from within the respective guest OS at a controlled rate, said controlled rate being adjusted up or down between a predetermined maximum rate and a predetermined minimum rate as the memory reservation software changes the allocation of the guest physical memory from within the respective guest OS,

said controlled rate being increased if said memory reservation software module is successful in changing the allocation of the guest physical memory from within the respective guest OS at the controlled rate, and

said controlled rate being decreased if said memory reservation software module is not successful in changing the allocation of the guest physical memory from within the respective guest OS at the controlled rate.



28. (Previously Presented): A system as in claim 27, in which the memory reservation software module is a driver installed within each respective guest OS.

29. (Previously Presented): A system as in claim 27, further comprising:  
a resource scheduler in the host system for allocating the hardware memory among the virtual computers;  
for each virtual computer, a communications module for communicating a respective memory quantity request to each driver,  
each driver being provided, upon sensing the respective memory quantity, for causing the corresponding guest OS to reserve an amount of the guest physical memory corresponding to the memory quantity request.

30. (Currently Amended): A computer system comprising:  
a host system, which includes at least one hardware processor, a host operating system (OS) and hardware memory that is addressable in a hardware memory address space; and  
a plurality of guest systems, each of which comprises a body of computer-executable code that is stored in the hardware memory and that addresses a guest memory in a guest address space, and is operatively connected to the host system;  
a memory reservation software module located within each guest system for receiving a memory quantity request from the host system and for changing the allocation of the guest memory from within the respective guest system according to the memory quantity request, thereby changing the amount of the hardware memory available for arbitrary use by the host system;  
in which:  
the guest address spaces of all the guest systems are mapped as portions of the hardware memory space;  
the guest systems are all executable on the same hardware processor(s), which are collocated in a single hardware platform,  
wherein the memory reservation software module, in response to the memory quantity request, changes the allocation of the guest physical memory from within the

respective guest OS at a controlled rate, said controlled rate being adjusted up or down between a predetermined maximum rate and a predetermined minimum rate as the memory reservation software changes the allocation of the guest physical memory from within the respective guest OS.

said controlled rate being increased if said memory reservation software module is successful in changing the allocation of the guest physical memory from within the respective guest OS at the controlled rate, and

said controlled rate being decreased if said memory reservation software module is not successful in changing the allocation of the guest physical memory from within the respective guest OS at the controlled rate.

31. (Previously Presented): A system as in claim 30, in which the memory reservation software module is a driver installed within each respective guest system.

32. (Previously Presented): A system as in claim 30, in which the memory reservation software module is a user-level application loaded in the guest system.

33. (Previously Presented): A system as in claim 30, further comprising:

a resource scheduler in the host system for allocating the hardware memory among the guest systems;

for each guest system, a communications module for communicating a respective memory quantity request to each memory reservation software module, each memory reservation software module being provided, upon sensing the respective memory quantity, for causing the corresponding guest system to reserve an amount of the guest memory corresponding to the memory quantity request.